



AWS PARTNER CERTIFICATION READINESS

Week 1 Content Review

Data Engineer – Associate



Today's Learning Outcomes



During this session, we will cover:

- AWS Global Infrastructure
- AWS Compute
- Networking and Security
- Storage and Databases





AWS PARTNER CERTIFICATION READINESS

AWS Global Infrastructure

Data Engineer - Associate

AWS Regions



A physical location around the world where AWS clusters data centers

What's in a Region?

Each AWS Region consists of multiple, isolated, and physically separate Availability Zones.

Why are they important?

AWS Regions are totally isolated from each other, creating the greatest possible fault tolerance and stability.



AWS Availability Zones (AZs)

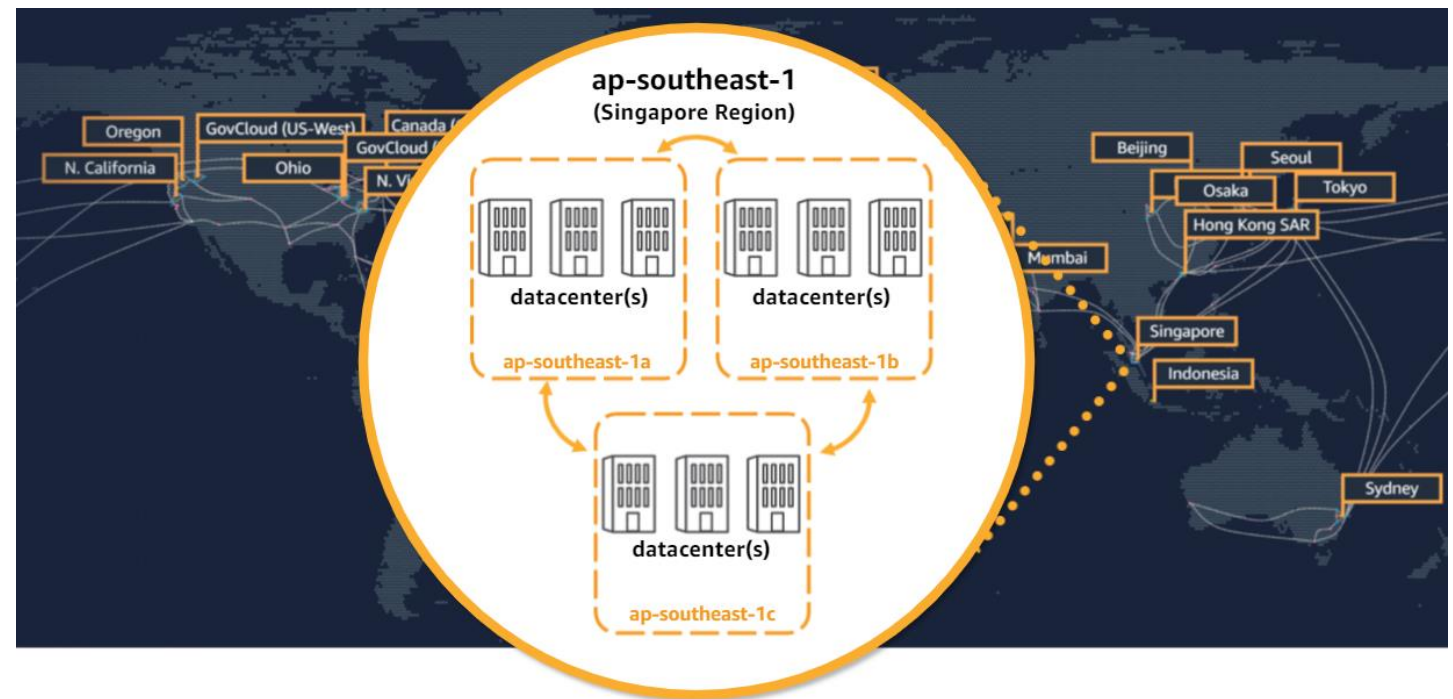


One or more discrete data centers with redundant power, networking, and connectivity in an AWS Region

Why are they important?

AZs give customers the ability to operate production applications and databases that are more highly available, fault tolerant, and scalable than would be possible from a single data center.

AZs are connected to each other with fast, private fiber-optic networking, enabling you to easily architect applications that automatically fail-over between AZs without interruption.



Edge Locations



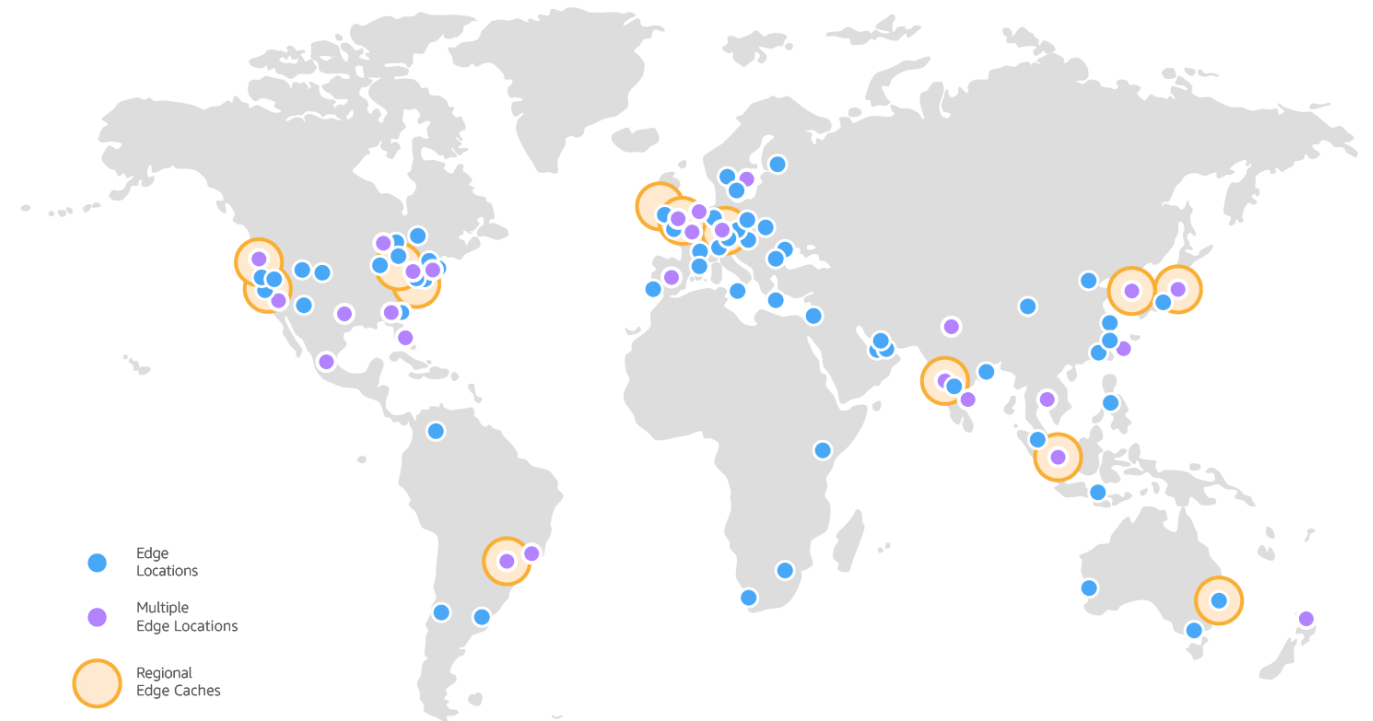
600+ CloudFront POPs and 13 regional mid-tier regional cache servers

What are they?

Smaller endpoints used for hosting cached data.

Why are they important?

Points of Presence (POP) enable Amazon CloudFront to securely deliver data, videos, applications, and APIs to customers globally with low latency and high transfer speeds, all within a developer-friendly environment.





AWS PARTNER CERTIFICATION READINESS

AWS Compute

Data Engineer - Associate



Amazon EC2



Secure and resizable compute capacity to support virtually any workload

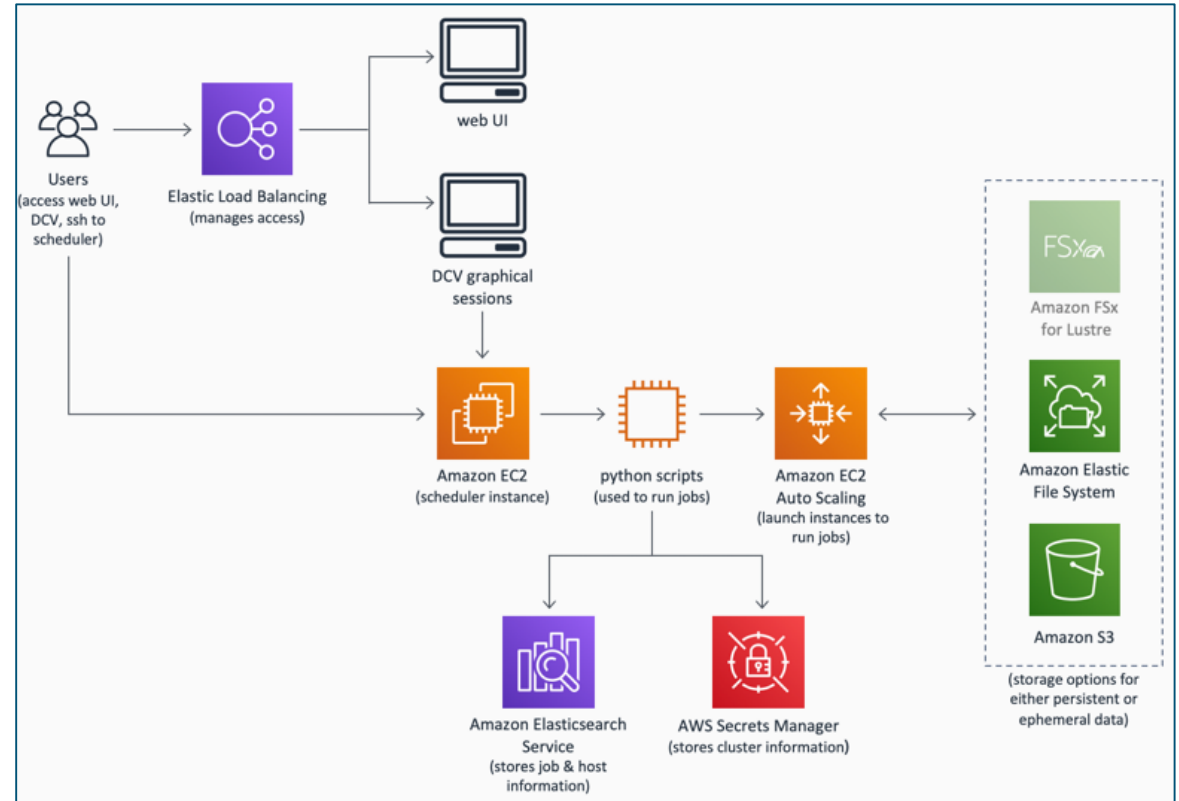
A web service that provides secure, resizable compute capacity in the cloud. It is designed to make web-scale cloud computing easier for developers

Scale

Increase or decrease capacity within minutes and provide 99.99% availability for each Amazon EC2 region

Save

Offers five pricing models to pay for Amazon EC2 instances: On-Demand, Savings Plans, Dedicated Hosts, Spot Instances and Reserved Instances.



Amazon Machine Images (AMI)



Critical information needed when launching EC2 instances

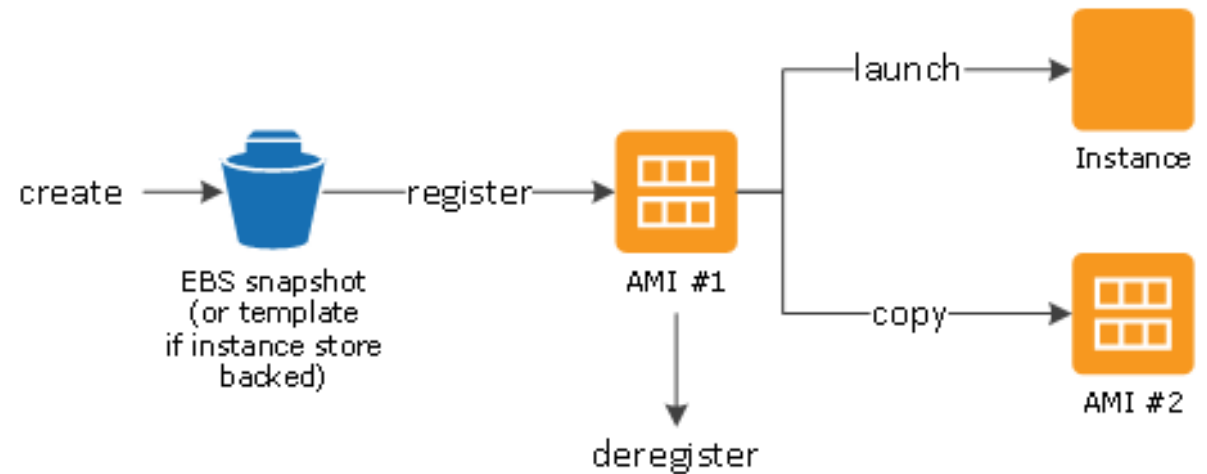
How it Works

Contains information required to launch an EC2 instance.

Must be specified when launching an instance.

Multiple instances can be launched from a single AMI.

Can include desired software to be installed on the EC2 instance at launch time.



What are Containers?



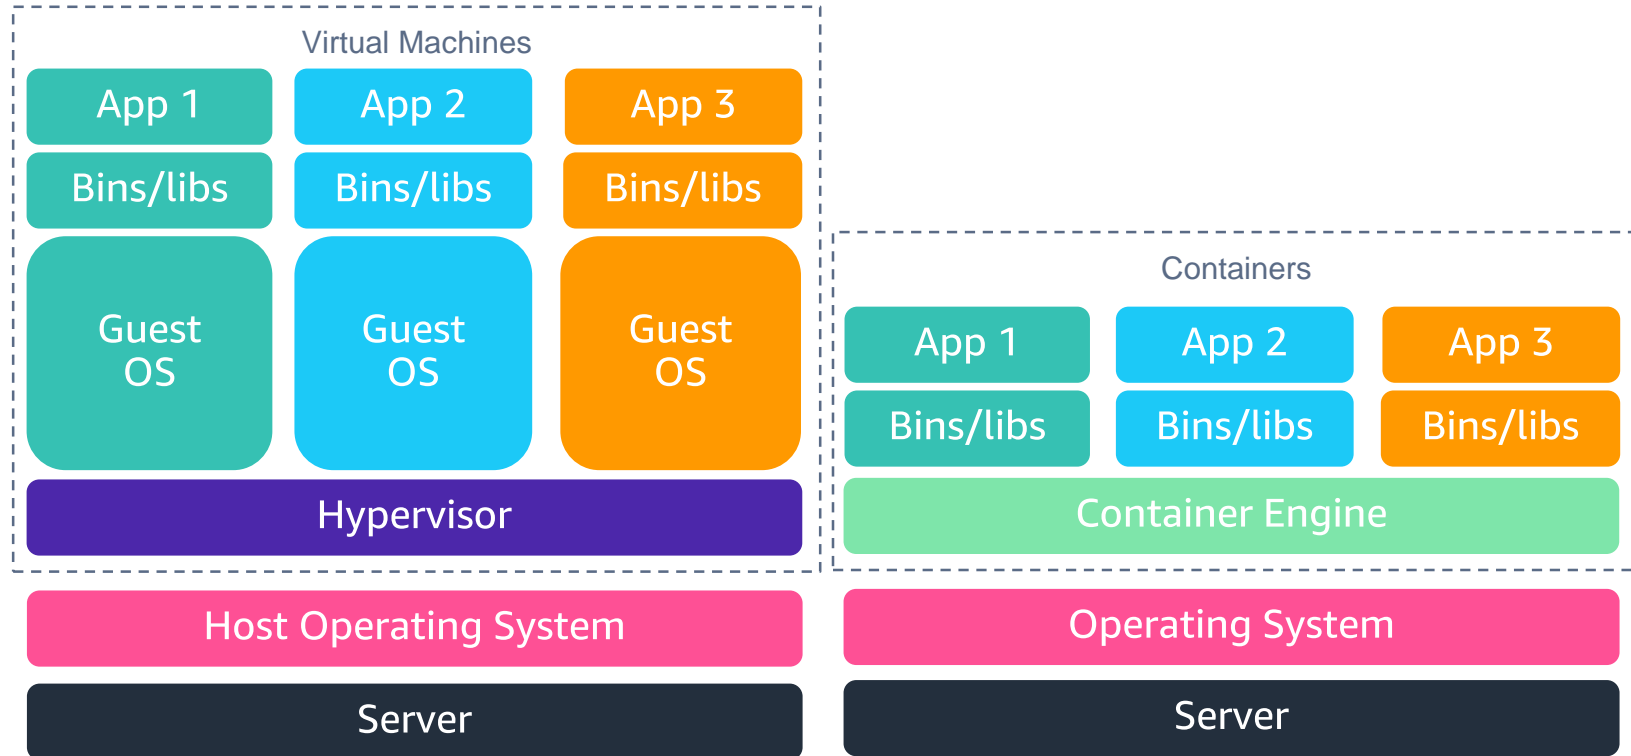
Lightweight, standardized packaging for application code, configs, and dependencies

Why

Portable method to package and deploy applications to run and scale anywhere. Suitable for deploying microservices, running batch jobs, for ML applications, and migrating applications to the cloud.

How it Works

Container images contain all the code, runtime, system libraries, dependencies, and configuration required for the application to run. Abstraction at the application layer allows containers to share the OS resources. Container engines "run" the images.



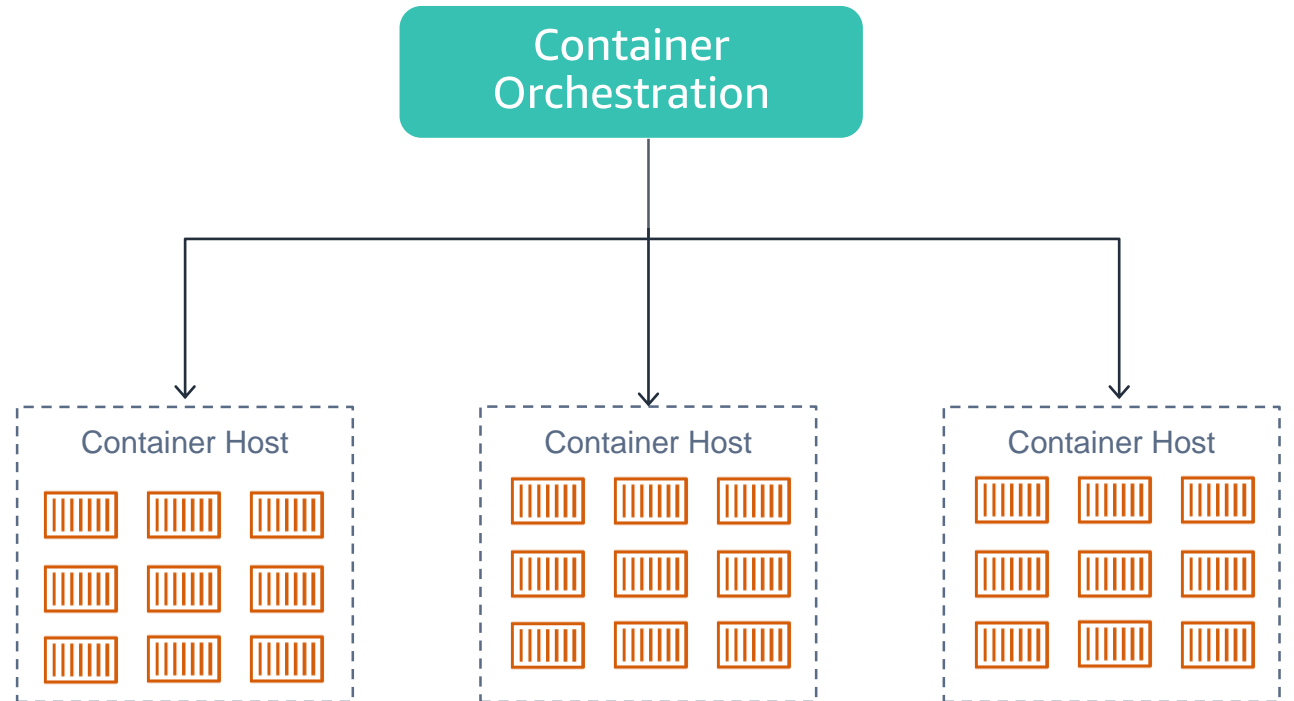
Container Orchestration



Automates scheduling, deployment, networking, scaling, health monitoring, and management of containers

What it is

Container orchestration automates the scheduling, development, networking, scaling, health monitoring, and management of your containers. Orchestration keeps containers running in the required state, and helps maintain your service-level agreements (SLAs).



Amazon Elastic Container Service



Highly secure, reliable, and scalable way to run containers

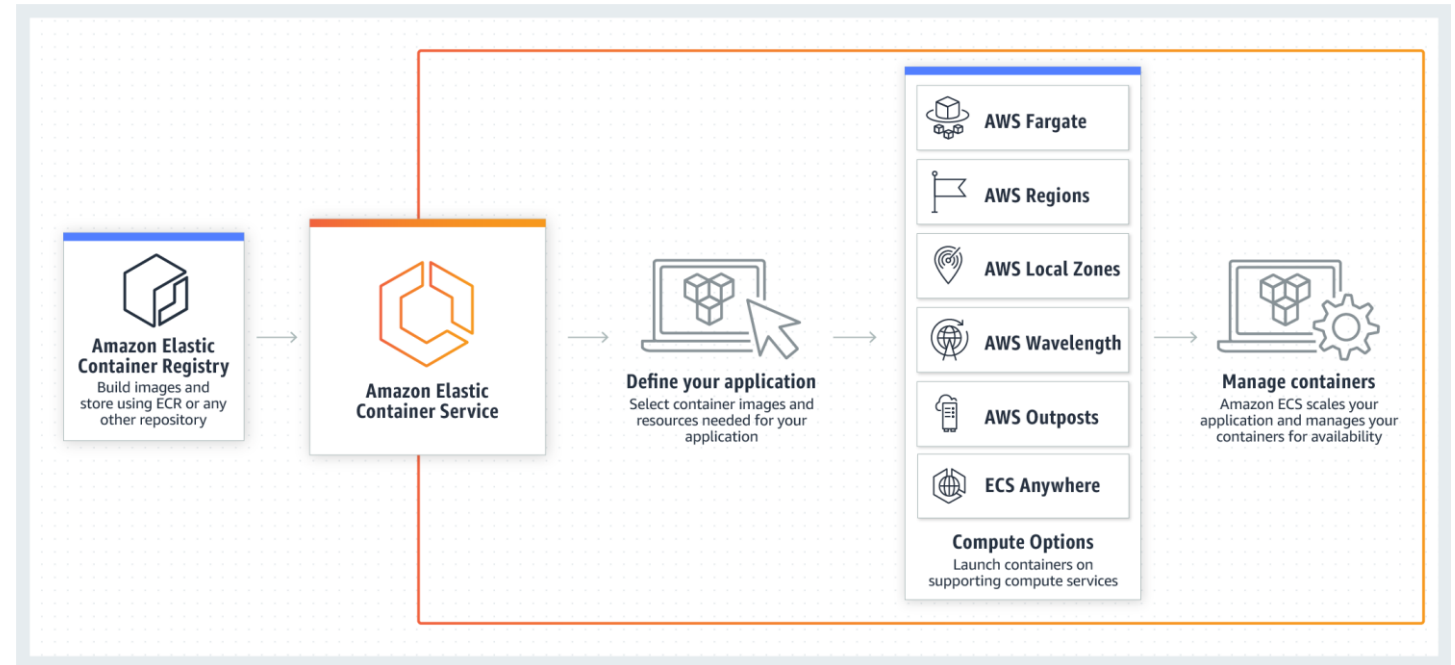
Fully managed container orchestration service that helps you easily deploy, manage, and scale containerized applications

Manage

Amazon ECS enables you to launch and stop your container-based applications by using simple API calls.

Offload

With Amazon ECS, you don't have to operate your own cluster management and configuration management systems or worry about scaling your management infrastructure.



Amazon Elastic Kubernetes Service



The most trusted way to start, run, and scale Kubernetes

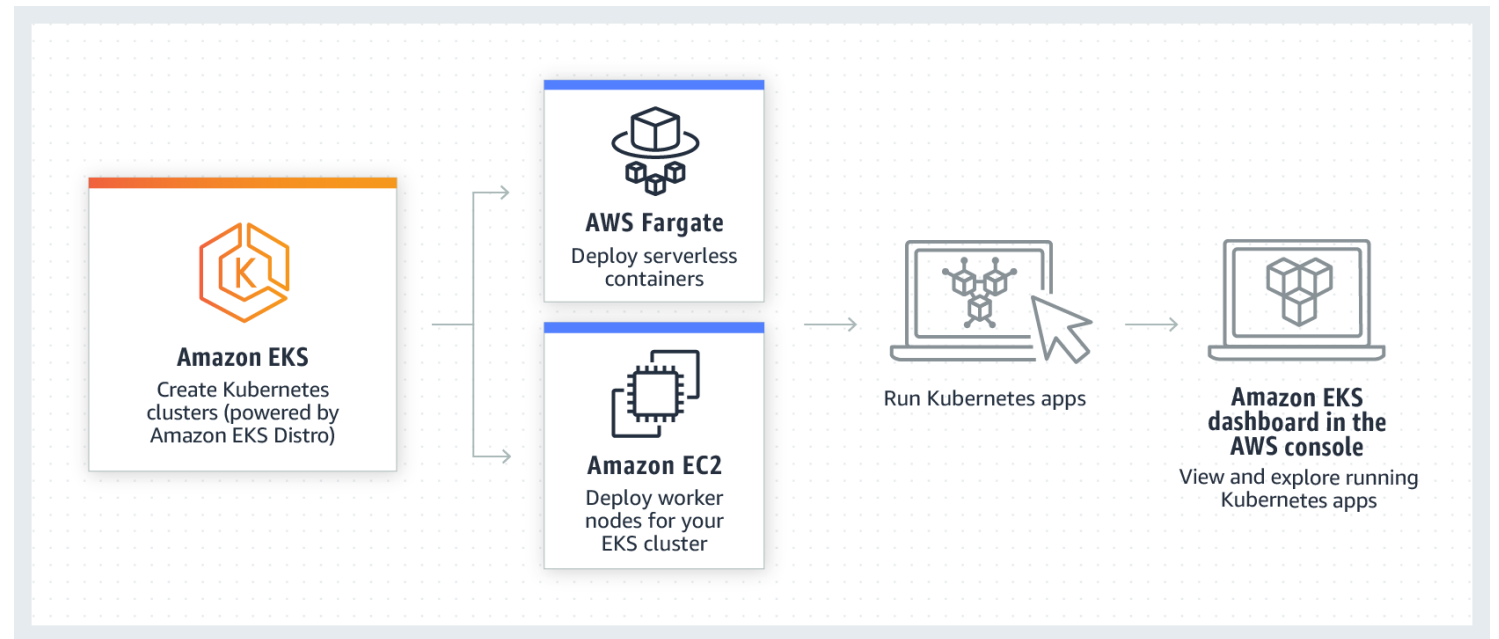
Gives you the flexibility to start, run, and scale Kubernetes applications in the AWS Cloud or on-premises. Runs upstream Kubernetes and is certified Kubernetes conformant

Availability

EKS runs the Kubernetes control plane across multiple Availability Zones, automatically detects and replaces unhealthy control plane nodes, and provides on-demand, zero downtime upgrades and patching.

Scalability

With EKS managed node groups, you don't need to separately provision compute capacity to scale your Kubernetes applications.



AWS Fargate



Serverless compute for containers

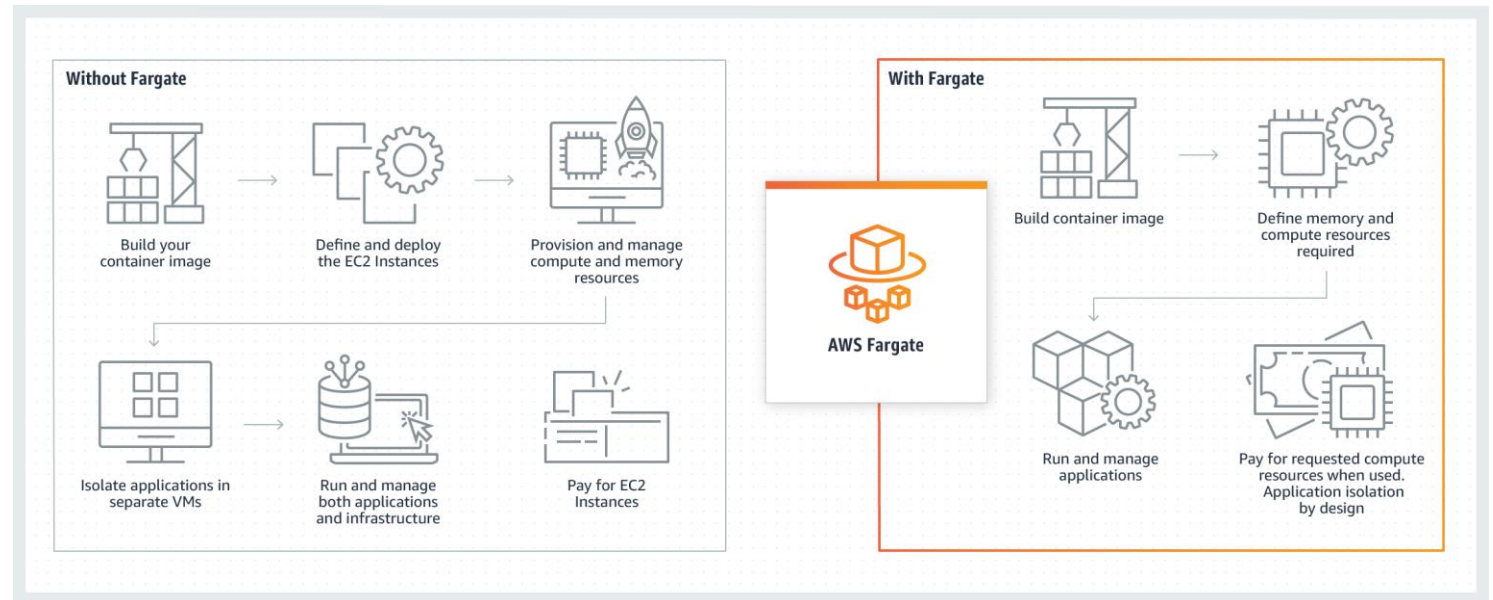
Serverless, pay-as-you-go compute engine that lets you focus on building applications without managing servers.

AI and ML Apps

Create a flexible and portable artificial intelligence (AI) and machine learning (ML) development environment.

Data Processing

Scale up to 16 vCPU and 120GB memory per task to run data processing workloads, with AWS Batch integration for serverless parallel processing.





AWS PARTNER CERTIFICATION READINESS

Scaling and Balancing

Data Engineer - Associate



Amazon EC2 Auto Scaling



Automatically launch or terminate Amazon EC2 instances


Fleet Management

Maintain health and availability of EC2 fleet. Monitor health of running instances, replace impaired instances automatically, and balance capacity across Availability Zones.

Scaling Options

Scheduled scaling allows you to scale up or down ahead of known load changes.

Dynamic scaling allows you to automatically follow the demand curve of your application usage based on load metrics.



Amazon EC2 Auto Scaling Groups




Collection of Amazon EC2 instances, treated as a logical grouping

What it Does

Allows usage of EC2 Auto Scaling features such as health check replacements, and scaling policies. Group launches enough instances to meet desired capacity and maintains by performing periodic health checks. The Auto Scaling Group is the defined group of instances which is managed by the EC2 Auto Scaling policies.

On-Demand instances, Spot instances, or both may be launched. Templates should be defined for the groups, with consideration of multiple Availability Zones.



Elastic Load Balancing



Distribute network traffic to improve the scalability of your applications

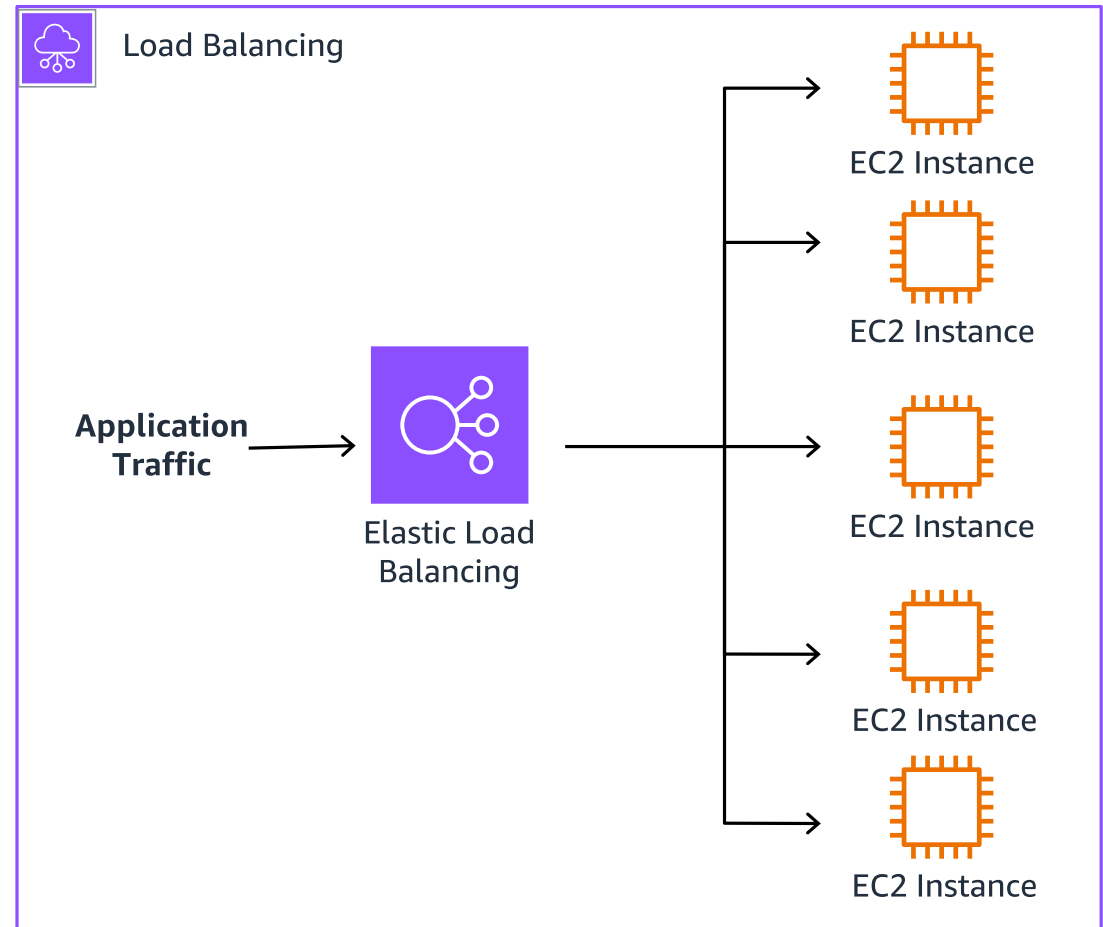
Automatically distributes incoming application traffic across multiple targets

Availability

Elastic Load Balancing is part of the AWS network, with native awareness of failure boundaries like AZs to keep your applications available across AZ's within a region

Single point of contact

Serves as the single point of contact for clients. The load balancer distributes incoming application traffic across multiple targets, such as EC2 instances, in multiple Availability Zones.





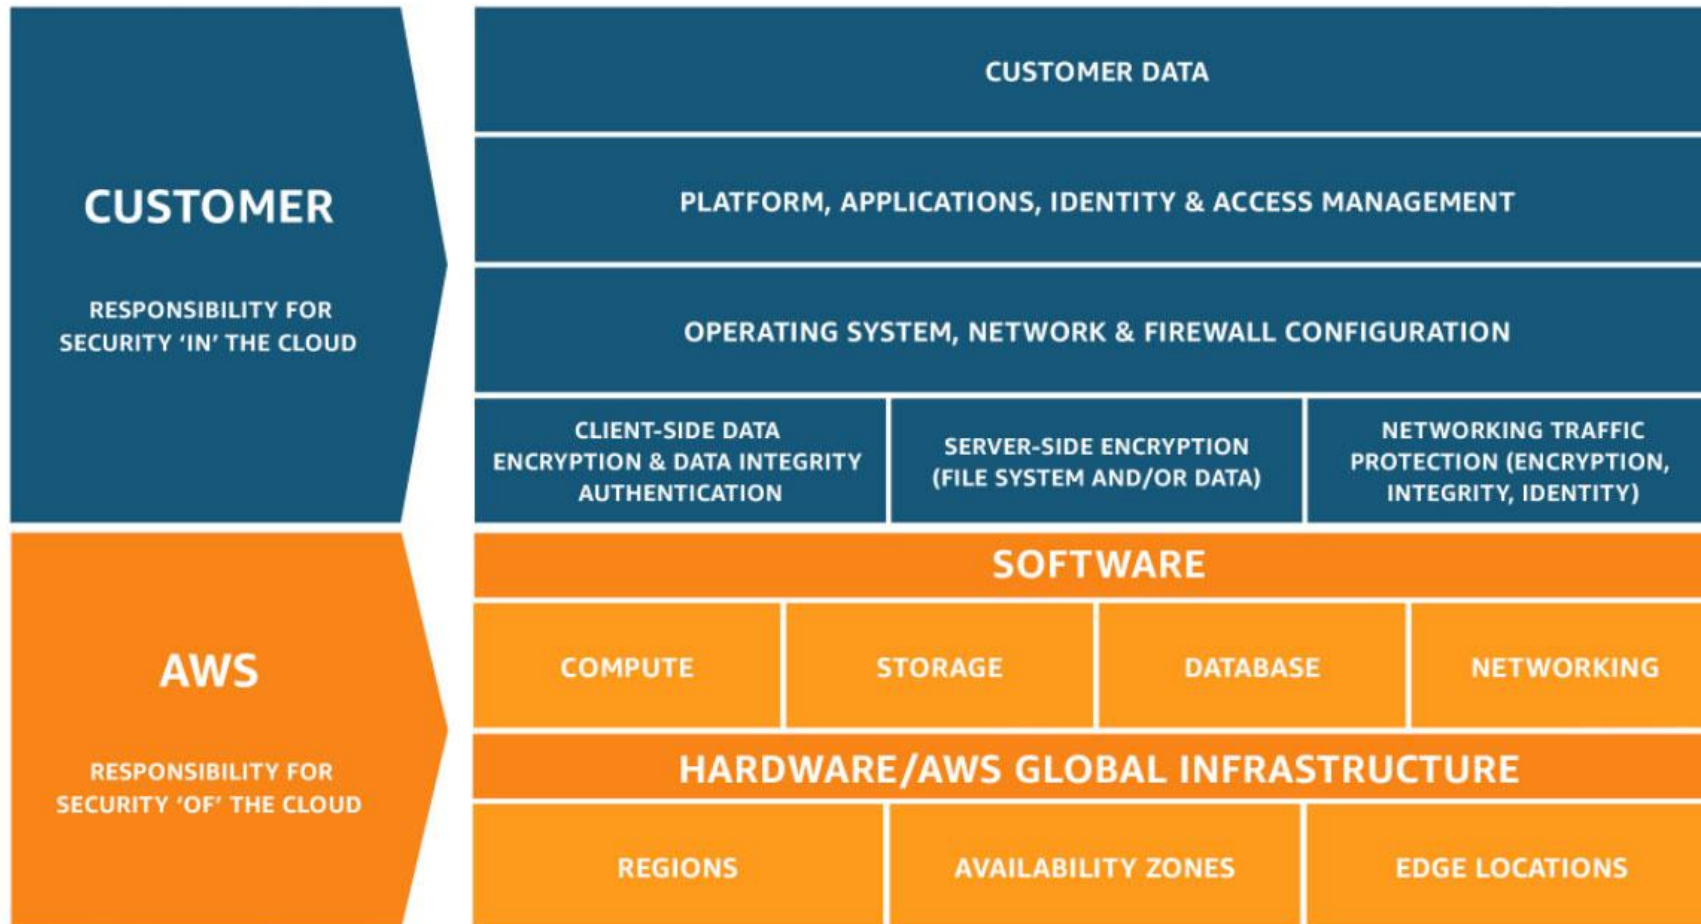
AWS PARTNER CERTIFICATION READINESS

Networking & Security

Data Engineer - Associate

AWS Shared Responsibility Model

Must have knowledge for Cloud Professionals



Amazon Virtual Private Cloud (VPC)



Provision a Logically Isolated Section of the AWS Cloud

Subnets

A range of IP addresses in your VPC. You can create AWS resources, such as EC2 instances, in specific subnets.

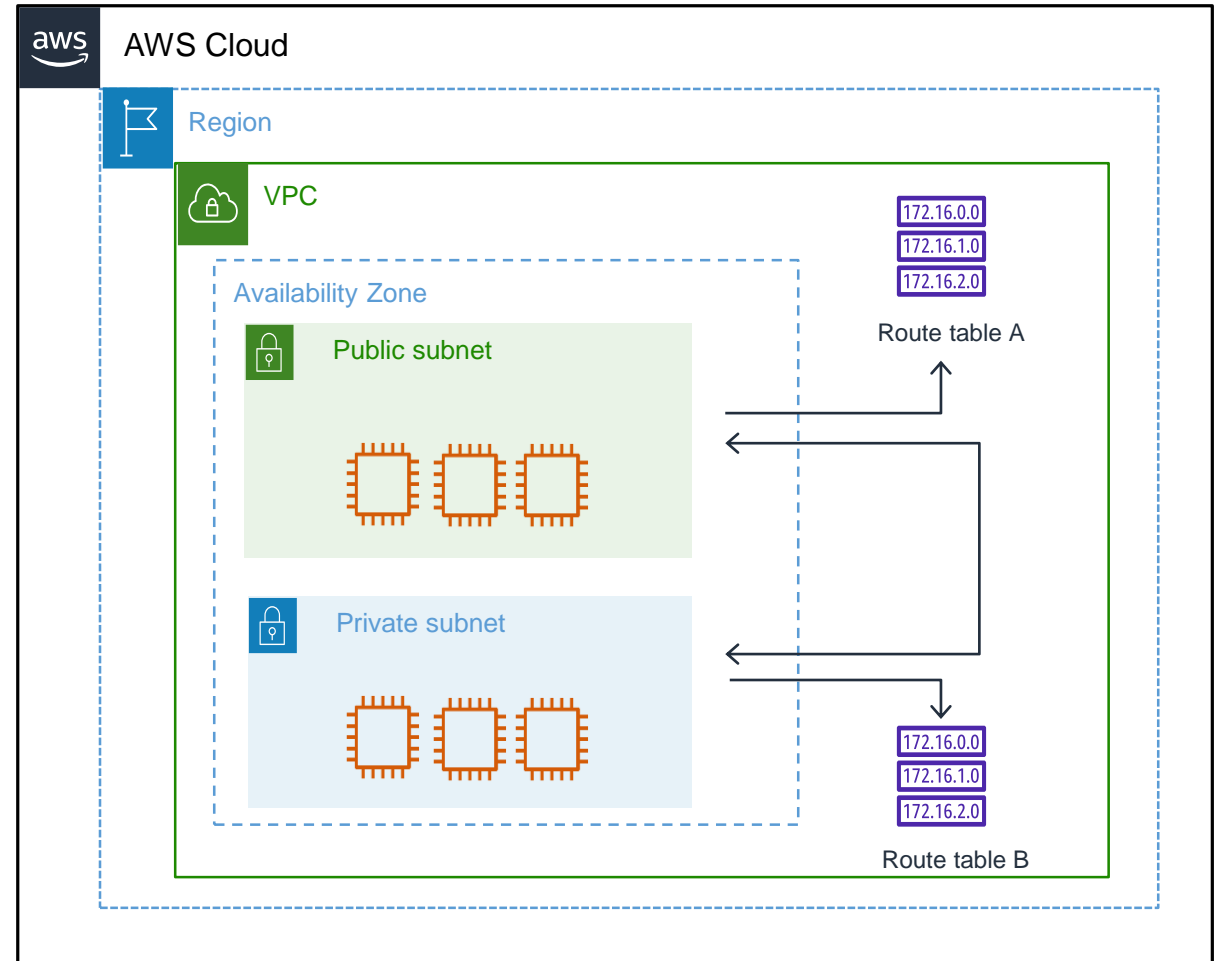
Can be **Public** or **Private**.

Cannot span availability zones.

Route Tables

Contains a set of rules, called routes, that determine where network traffic from your subnet or gateway is directed.

You use route tables to control where network traffic is directed. Each subnet in your VPC must be associated with a route table, which controls the routing for the subnet (subnet route table).



Security Groups & NACLs

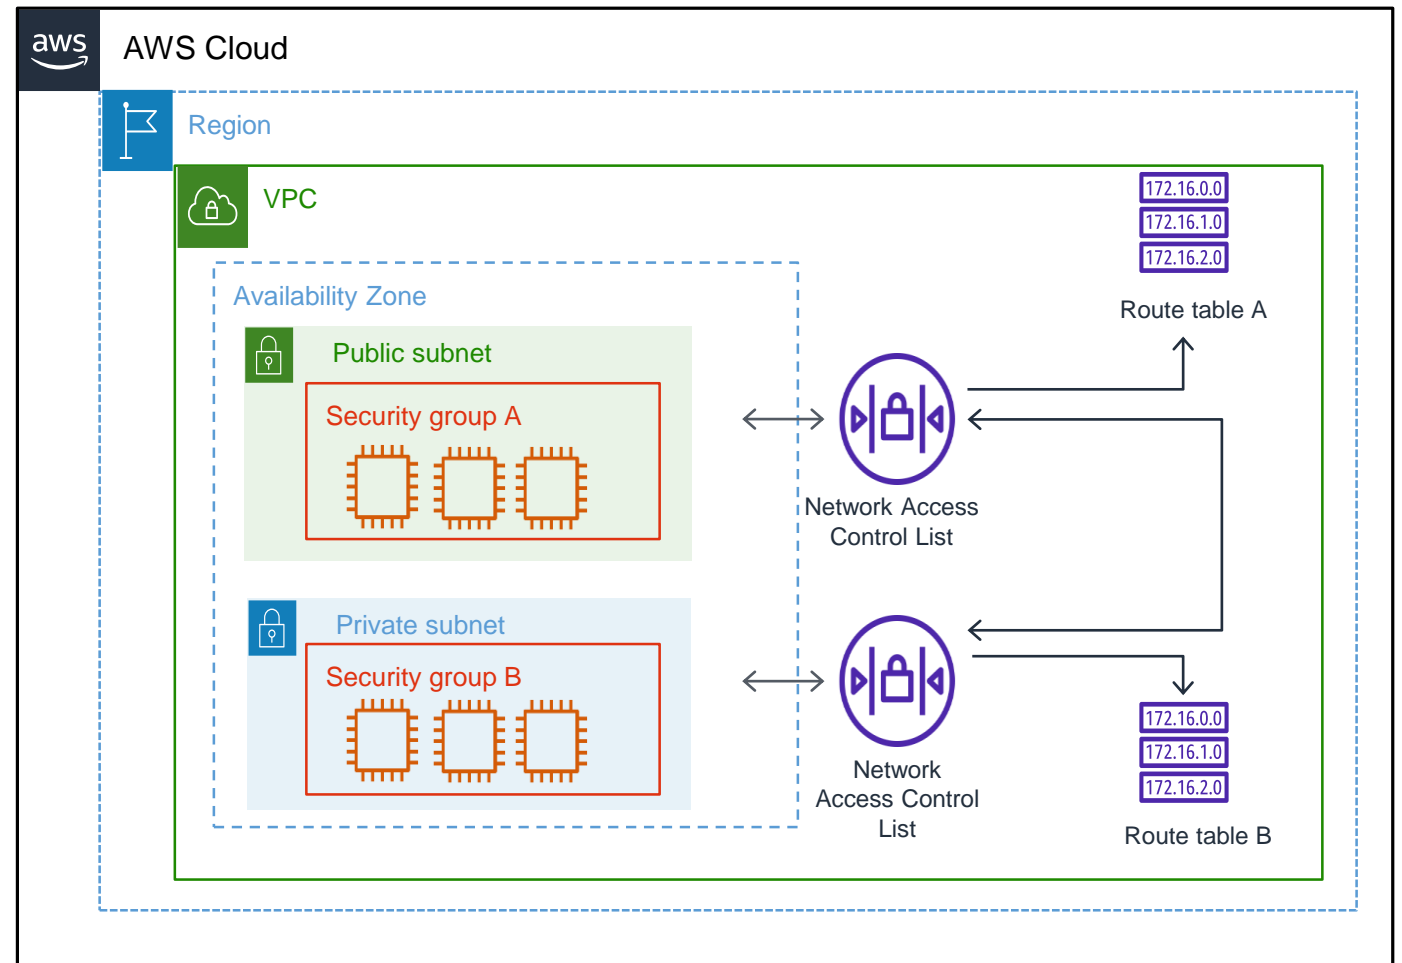
Layered security options for your VPC – virtual firewalls

Security Groups

Act as a firewall for associated **Amazon EC2 instances**, controlling both inbound and outbound traffic at the instance level.

Network Access Control Lists

Act as a firewall for associated **subnets**, controlling both inbound and outbound traffic at the subnet level.



Amazon Identity and Access Management (IAM)



Enables AWS customers to manage access to AWS services and resources

Provides recommendations that help you follow AWS best practices.

How does it work?

Trusted Advisor can help you save cost, such as recommending you to delete unused or idle resources, or use reserved capacity.

IAM Identities


Users - identity that has specific permissions for a single person or application


User Groups - identity that specifies a collection of IAM users

Roles - identity with specific permissions. Assumable, not associated with specific person.

 Fine-grained access control to AWS resources

 Multi-factor authentication

 The ability to analyze access

 Integration with corporate directories

Root User Accounts



First AWS account, which has complete access to all AWS services and resources

Important!

- **DO** – Enable Multi-factor authentication (MFA) on this account
- **DO** – set a secure password.
- **DO** – Use this account to create an Admin IAM user
- **DO NOT** – Use this account for every-day tasks
- **DO NOT** – Create access keys for the root user unless absolutely necessary





AWS PARTNER CERTIFICATION READINESS

Storage and Databases

Data Engineer - Associate



Managing Data Within AWS



Data Storage



EFS



FSx

File



EBS

Block



S3



S3 Glacier

Object

What is File Data?

A hierarchical storage system that provides shared access to file data.

Files also contain metadata like file name, size, and timestamps

What is Block Data?

Works by dividing data into fixed-sized blocks and stores them as individual units.

Includes most data types

What is Object Data?

Stores and Manages data as Objects

Examples include data like documents, images, or data values

Amazon MemoryDB for Redis



Redis-compatible, durable, in-memory database service for ultra-fast performance

Amazon MemoryDB for Redis is a durable database with microsecond reads, low single-digit millisecond writes, scalability, and enterprise security

Use Cases

- Workload requires ultra-fast performance
- Building application with Redis data structures and APIs with a primary, durable database
- Simplifying application architecture by replacing usage of a database with a cache



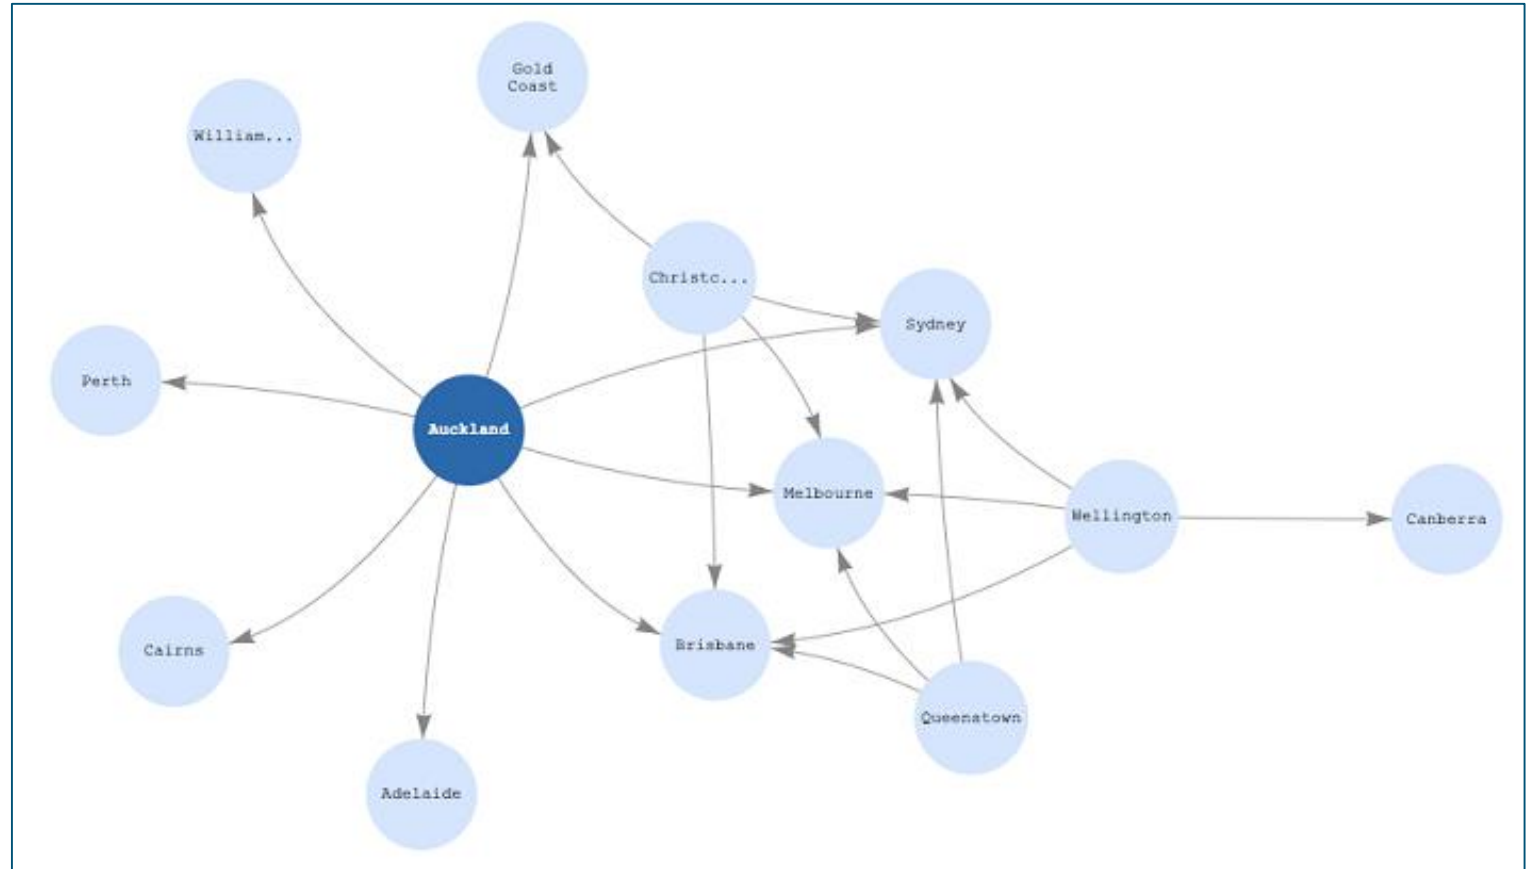
Amazon Neptune



High performance graph analytics and serverless database for superior scalability and availability

Database

Amazon Neptune Database provides a purpose-built graph database with a full set of enterprise features and integrations. Neptune Database supports mission-critical graph applications that require high availability, disaster recovery, dynamic scaling, and other capabilities required by enterprise applications.



Amazon Keyspaces (for Apache Cassandra)




A scalable, highly available, and managed Apache Cassandra-compatible database service

Cassandra Compatibility

Compatible with the open-source Cassandra Query Language (CQL) API, so you can migrate your existing Cassandra tables to Amazon Keyspaces while continuing to use your existing application code.

High Speed Data Processing

Process data at high speeds for applications that require single-digit-millisecond latency, such as industrial equipment maintenance, trade monitoring, fleet management, and route optimization.



Amazon DocumentDB (MongoDB Compatible)



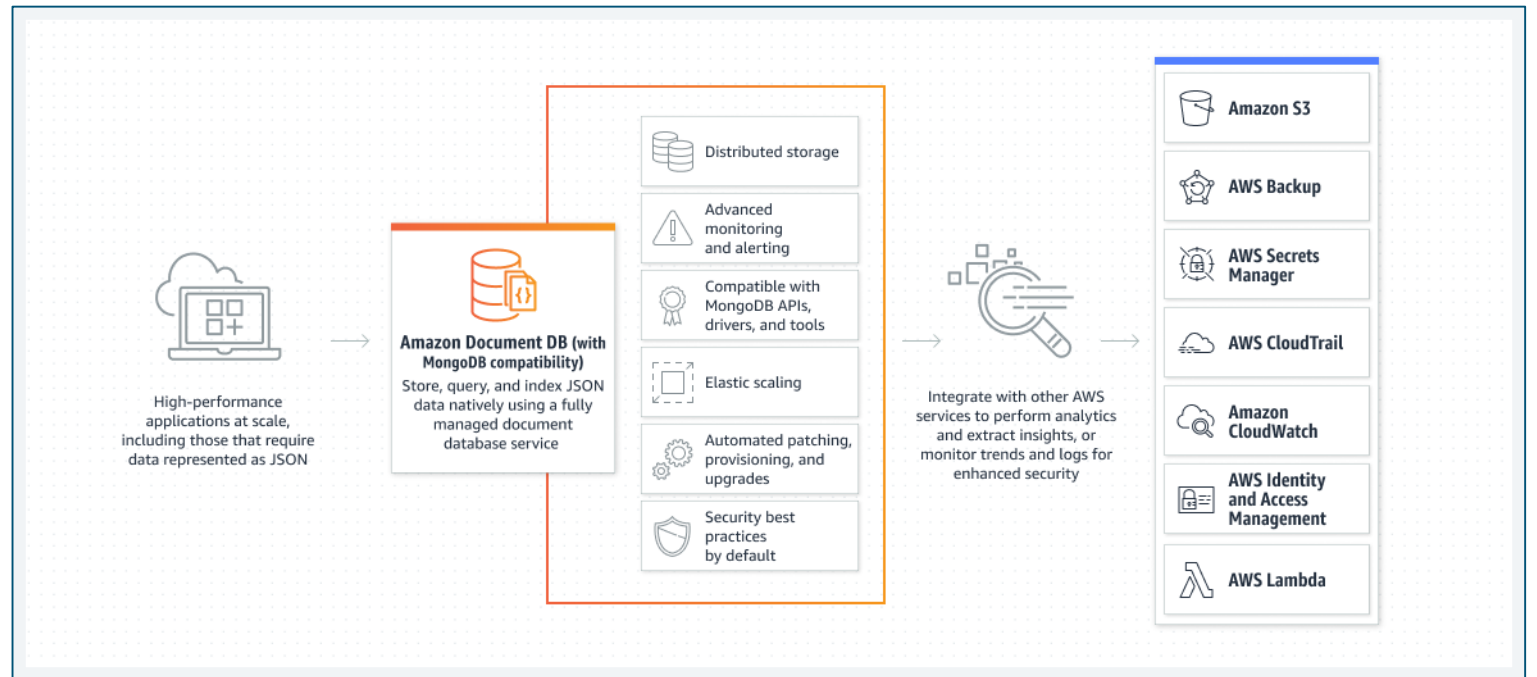
Scale enterprise workloads with ease using a fully managed native JSON document database

Content Management Data

Improve the customer experience through fast, reliable access to reviews, images, and other content stored in your content management system (CMS).

Enhance Applications

Apply generative AI and ML capabilities to unlock a wide range of use cases including semantic search experiences, product recommendations, personalization, chatbots, and more.





**Thank you for attending
this session**